FIELD WORKBOOK

Is Your LBM Business AI-Ready?

A field workbook for Epicor Insights for Building Supply

Kristin Andrews · Director of Process Optimization, MANS Lumber & Home · kandrews@manslumber.com

How to use this workbook

Grab a pen, or type right in it, and answer honestly. The answers you'd rather skip are the ones that matter most.

Who it's for: owners, managers, and crew at building-supply and LBM businesses. The folks at Epicor Insights who run BisTrack every day.

What "AI" means here: practical tools you can use Monday morning. ChatGPT, Claude, Microsoft Copilot, and the AI features showing up inside the systems you already run. Not robots, and not sci-fi.

How to work through it:

Section 1	Score yourself honestly
Section 2	Fix the foundations before you buy anything
Section 3	Pick one first move
Section 4	Bring your crew with you
Section 5	Learn the one skill that makes AI useful
Section 6	Set guardrails so you don't get burned

> **One rule:** don't skip to the shiny part. Most AI projects fail on the foundations, not the technology.

Section 1: Is your business actually ready?

Score each statement: **2** = true, **1** = sort of, **0** = no.

#	Statement	Score (0–2)
1	Our BisTrack data is clean enough to trust. Customers, pricing, and inventory aren't a mess.	
2	There's a specific problem we want to solve, beyond "we should do AI."	
3	One named person owns this. It isn't "everybody's" job.	
4	Leadership backs it with budget and air cover.	
5	We can free up a little time to learn and test.	
6	At least one person on the team is curious about it.	
7	We have basic security habits: passwords, who sees what.	
8	We can say how we'll measure success before we start.	
9	We understand AI assists people. It doesn't run unattended.	
10	Our expectations are realistic. We want help, and we're not expecting magic.	
	TOTAL	

Your verdict

> **16 to 20, Ready.** Go pick a pilot in Section 3.

> **10 to 15, Almost.** Work on your two or three lowest scores first (Section 2), then start small.

> **Under 10, Not yet.** Don't buy anything. Fixing the foundations costs far less than a failed rollout.

Honesty check: if you gave yourself a 2 on everything, score it again and be tougher. A test nobody can fail tells you nothing.

Section 2: Before you decide anything

Skip these five and any tool will flop.

1. Clean data beats smart AI. AI working from messy BisTrack data gives confident, wrong answers. Fix the worst problems first: duplicate customers, stale pricing, junk SKUs.

2. Name one owner. One person drives it, even if the project is small.

3. Point at a real problem. Write down one specific pain, like "quotes take two days to follow up." "Be more efficient" doesn't count.

4. Set a budget and a time box. Decide up front what you'll spend and how long you'll test before you judge it. For a first pilot, 30 to 60 days is plenty.

5. Pick your baseline number now. You can't prove it worked if you never measured "before."

Foundation checklist

- ☐ Worst data problems identified and being cleaned
- ☐ One owner named: _____
- ☐ The one problem, in writing: _____
- ☐ Budget set and test window agreed
- ☐ Baseline metric and today's number written down: _____

Section 3: Pick your first move

Your first project needs to be a win. Run every idea through three filters:

High pain	It annoys people or costs real money today.
Low risk	A mistake is cheap and gets caught before a customer sees it.
Measurable	You can count the before and after.

For your first try, stay away from anything customer-facing or anything that commits money.

Good first moves in an LBM shop

- Drafting quote follow-up emails from BisTrack order data
- Exporting and summarizing long BisTrack reports or smart views into a plain English page
- Analyzing reports or smart views to look for trends or gaps
- Writing first-draft job descriptions and help wanted ads
- Turning product spec sheets into quick customer answers
- Cleaning up and categorizing messy vendor data
- First-draft SOPs for yard and delivery processes

Rate your candidates (score each 1 to 5)

Candidate task	Pain	Low risk	Measurable	Total

Highest total wins. Finish that one project before you start another.

Section 4: Getting your crew on board

The technology is the easy part. Rollouts succeed or fail on people.

> The fear is specific: "This is here to replace me." Address it early, out loud, and mean it. The nail gun didn't replace framers. It let them build more.

Five moves that work

1. Involve them before you buy. People support what they help build.

2. Start with the task they hate most, so the first thing they see is a win.

3. Show them. A 10-minute live demo beats a policy memo.

4. Name a champion. One respected coworker using it well wins over more people than any manager can.

5. Celebrate the first win out loud. "Dave cut his quote follow-up time in half" spreads faster than any mandate.

> **What kills a rollout:** forcing it from the top, hiding how it works, or letting it feel like surveillance.

Worksheet: three tasks your team openly hates

1. _____

2. _____

3. _____

Start your pilot on the one that gets the most groans.

Section 5: Writing a prompt that actually works

A vague prompt gets a vague answer. The fix is a simple four-part framework, **R-T-C-F**:

Role	who the AI should act as
Task	what you want done
Context	the details only you know
Format	what the output should look like

Before (what most people type)

> *write a quote follow up email*

After (R-T-C-F)

> **Role:** You're an inside sales rep at a building-supply dealer.
>
> **Task:** Write a short, friendly follow-up on an open quote.
>
> **Context:** Customer is Miller Construction, quote #4821 for framing lumber, sent 5 days ago, no reply. We can hold the price through Friday.
>
> **Format:** Under 120 words, warm but direct, with one clear ask: confirm the order or call me.

Same tool, but the second prompt gets you an email you can actually send.

Another example: sales manager or owner

> **Role:** You're a sales manager at a building-supply dealer who coaches reps on growing their accounts.
>
> **Task:** Review this rep's sales data, tell me what's working and what isn't, and give me three specific recommendations to improve their numbers.
>
> **Context:** Attached is a BisTrack smart view export for one rep: the last 12 months of sales, gross margin %, and quote-to-order rate by customer. Our target margin is 28%. Customer names are replaced with account numbers.
>
> **Format:** A short summary of the trends first, then three recommendations as bullets, each tied to the customers or numbers behind it. Keep it plain enough to walk through with the rep in a one-on-one.

Then push back: ask what it would do for the sales team as a whole, or what you might be missing. Only run this kind of analysis in a business version that doesn't train on your data (see Section 6).

Your fill-in template

Role: You are

Task:

Context:

Format:

A few more tips

- If the first answer isn't right, tell it what to fix. It's a conversation, not a vending machine.
- Give it an example of a good result when you have one. Examples work better than instructions.
- **ALWAYS push back.** AI is programmed to give the most likely next answer. This means that over time it becomes a reflection of you. It is important to always ask things like “Are we missing anything?”, “Is there anything that could make this better?”, or “Have we looked at this from the customer’s point of view?” By doing this, it forces the AI to look at different answers and not just tell you what it thinks you want to hear.

Section 6: AI safety and guardrails

Treat AI like any sharp tool: useful, and worth respecting.

1. Protect your data. Free public AI tools may train on what you paste, so treat anything you enter as potentially public.

> **Never paste these into a public tool:**
>
> - Customer personal info (addresses, phone or email lists)
> - Negotiated pricing, margins, or vendor contracts
> - Employee records, SSNs, payroll
> - Card or bank numbers
> - Anything you wouldn't post on your front door

Use the paid business or enterprise versions, which by contract don't train on your data, and set one simple company rule for what's allowed.

2. Keep a person in the loop. AI writes the draft and a person approves it. This is non-negotiable for anything with a price, a promise, or a customer's name on it.

3. Don't trust any number it gives you. AI will get prices, SKUs, measurements, and codes wrong and sound completely sure. Check every figure and spec before it leaves the building.

4. Watch for confident nonsense. "Hallucination" means the AI made something up. If an answer sounds too tidy, check the source.

Your guardrail checklist

- ☐ We use a version that doesn't train on our data
- ☐ We have a written "never paste this" list
- ☐ A person approves anything going to a customer or involving money
- ☐ We spot-check facts and numbers
- ☐ One person owns our AI rules and reviews them every quarter

Pick one task, run it for 30 days, and measure it. Nothing beats seeing it work in your own business.

Questions? Let's talk. Kristin Andrews · Director of Process Optimization, MANS Lumber & Home · kandrews@manslumber.com